arXiv:1511.02210v1 [cs.AI] 6 Nov 2015

Learning Optimized Or's of And's

Tong Wang
CSAIL and Sloan, MIT
tongwang@mit.edu

Cynthia Rudin
CSAIL and Sloan, MIT
rudin@mit.edu

Abstract

Or's of *And*'s (OA) models are comprised of a small number of disjunctions of conjunctions, also called disjunctive normal form. An example of an OA model is as follows: If ($x_1 =$ 'blue' *AND* $x_2 =$ 'middle') *OR* ($x_1 =$ 'yellow'), then predict $Y = 1$, else predict $Y = 0$. Or's of And's models have the advantage of being interpretable to human experts, since they are a set of conditions that concisely capture the characteristics of a specific subset of data. We present two optimization-based machine learning frameworks for constructing OA models, Optimized OA (OOA) and its faster version, Optimized OA with Approximations (OOAx). We prove theoretical bounds on the properties of patterns in an OA model. We build OA models as a diagnostic screening tool for obstructive sleep apnea, that achieves high accuracy with a substantial gain in interpretability over other methods.

1 Introduction

We present mathematical programming formulations for producing *Or*'s of *And*'s, which are sparse disjunctive normal form (DNF) expressions. An OA model might say, for instance, consumers who are female *AND* single, *AND* younger than 35 years old, *OR* married *AND* earn more than \$100K per year, are likely to purchase a product. In creating predictive models for healthcare, marketing, sociology, and in other domains, two aspects have long since been of interest: logical forms, and sparsity (see, e.g. Dawes, 1979; Johnson-Laird, Khemlani, and Goodwin, 2015; Miller, 1956). For example, physicians use sparse, easily checkable sets of conditions (symptoms, observations) to classify patients as to whether they have a disease. DNF formulae are particularly useful as *screening* models, where patients who do not meet the Or of And's criteria are not considered for further testing. In marketing, DNF is also called "disjunctions of conjunctions" or "non-compensatory decision rules." Marketing researchers strongly hypothesize that consumers use simple rules to screen products, and they would consider purchasing only the products in this *consideration set* to reduce the cognitive load of considering all products. The consideration set may be precisely an Or's of And's classifier (Hauser et al., 2010; Gilbride and Allenby, 2004).

Despite the efforts that theoretical communities have placed on learning DNF (e.g., Klivans and Servedio, 2001; Littlestone, 1988; Ehrenfeucht et al., 1989), the algorithms designed for inductive logic programming that essentially produce DNF (Muggleton and De Raedt, 1994; Lavrac and Dzeroski, 1994), the associative classification algorithms (Ma, Liu, and Hsu, 1998; Han, Pei, and Yin, 2000; Li, Han, and Pei, 2001; Yin and Han, 2003; Chen et al., 2006; Cheng et al., 2007), and rule induction methods (e.g., Cohen, 1995), there has been little in the way of algorithms designed for applications where cognitive simplicity is first and foremost (with some exceptions, like Hauser et al., 2010, which we discuss later). For instance, Yin and Han (2003) reported that the average number of rules used by CMAR was 305 and CPAR had 244 rules on average, on 26 datasets from UCI ML Repository (Lichman, 2013); these are not cognitively simple models. Besides, all of the algorithms discussed above use greedy approximations, which hurts accuracy and sparsity. For instance, RIPPER employs local greedy splitting, meaning that a mistake at the beginning is difficult to undo. Inductive logic programming starts with a collection of rules and locally combines them. The associative classification methods also follow separate-and-conquer or covering strategies. Unlike these methods, our methods aim to produce cognitive simple models, and do not use greedy approximations or similar heuristics.

The closest work to ours are that of Hauser et al. (2010) in the marketing literature, and Wang et al. (2015) on Bayesian modeling of DNF formulae. Hauser et al. (2010) pre-mines the rules and then uses an integer program equivalent to a set-covering problem. They try to minimize the

error of coverage while favoring patterns that fit the largest subset of data. The work of Wang et al. (2015) has the advantage of a Bayesian interpretation of the prior parameters, but a disadvantage in that the analytical bounds on the maximum a posteriori solution of the Bayesian method are weaker than those we present in this paper for the minimizer of the optimization problem.

Our goal is cognitive simplicity, as well as predictive accuracy. We choose mathematical programming (mixed-integer and integer linear programming – MIP and ILP) and rule mining to form our models. Using these tools have several benefits, namely flexibility on the user's side on the objective and constraints, fast solvers that have been improving exponentially over recent years, and a guarantee on the optimality of the solution. We improve computation also using statistical approximations. In one of our algorithms, OOAx, we first mine rules and design an ILP to choose the subset of rules to form the OA model. This is a statistical assumption that dramatically speeds up computation, but we can show (in Theorem 2) that as long as we mine all rules with sufficiently high support, the optimal solution will be attained anyway. We also present various bounding conditions on the optimal solution.

2 Optimized Or's of And's

Let us discuss the first framework for learning Or's of And's, called Optimized Or's of And's (OOA). We work with a data set $S = \{(X_n, Y_n)\}_n^N$, consisting of N examples with J attributes of mixed type. $Y_n \in \{1, -1\}$ represents the labels. Numerical attributes are indexed by index set $\mathcal{J}_{\text{n}}$ and categorical attributes are indexed by $\mathcal{J}_{\text{c}}$. The j-th attribute of the n-th example is denoted as X_{nj}.

An OA classifier consists of a set of patterns that characterize a single class, here, the positive class. Each pattern is a conjunction of conditions (literals), and the number of conditions is called the *length* of a pattern. For example. the length of pattern "age $\geq$ 30 *AND* has hypertension *AND* is female" is 3. Let z denote a pattern, and $\mathbb{1}_z(X)$ indicate if X satisfies pattern z. A represents a set of patterns. An OA classifier built on A is denoted as f_A:

$$f_A(X) = \begin{cases} 1 & \exists z \in A, \text{s.t.} \mathbb{1}_z(X) = 1 \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

2.1 MIP Formulation

We formulate a mixed integer program to generate a pattern set containing numerical and categorical attributes. The MIP uses the following objective $L(A)$ to minimize the training error while maintaining sparseness of the model.

$$L(A) = \tfrac{\#\text{error}(A)}{N} + C_1 \#\text{literals}(A) + C_2 \#\text{patterns}(A). \quad (2)$$

The first term in the objective is the loss function, which counts the number of classification errors. The regularization terms include 1) the total number of literals in A, denoted as $\#\text{literals}(A)$, which is the sum of the length of each pattern in A, and 2) the total number of patterns in A, denoted as $\#\text{patterns}(A)$. The two terms are scaled by parameters C_1 and C_2 to penalize the complexity of the model. C_1 represents the percentage of training errors the user is willing to trade in order to reduce a pattern by one literal. Similarly, C_2 represents the percentage of training errors a user needs to trade to reduce one pattern. A user can tune C_1 and C_2 to influence the shape of the output.

Now we explain how the constraints work. A challenging part is to deal with both numerical and categorical attributes in the MIP. For numerical attributes, the MIP needs to select the upper and lower boundary to form a range; for categorical attributes, it needs to select a category for a literal. Simultaneously, the MIP needs to decide for each example if it satisfies the literals and patterns. All of the constraints are linear in the decision variable, to ensure a duality gap or proof of optimality on the solution we obtain.

2.1.1 Literals for Numerical Attributes

For a numerical attribute, a literal j (for simplicity, when we refer to literal j, we mean a literal containing attribute j) has the form "$l_{kj} \leq X_{\cdot j} \leq u_{kj}$", where u_{kj} and l_{kj} represent the upper and lower boundary of the range in literal j. k is a pattern index. For each example X_n, let $\hat{u}_{nkj} \in \{0, 1\}$ indicate if X_{nj} satisfies the upper bound u_{kj} and $\hat{l}_{nkj} \in \{0, 1\}$ indicate if X_{nj} satisfies the lower bound l_{kj}. That is, $\hat{u}_{nkj} = 1$ if $X_{nj} \leq u_{kj}$, and $\hat{l}_{nkj} = 1$ if $X_{nj} \geq l_{kj}$. Using a big M formulation, we obtain the following constraints, that for $\forall n, k, \forall j \in \mathcal{J}_{\text{n}}$,

$$u_{kj} - X_{nj} \leq M\hat{u}_{nkj}, \quad (3)$$
$$u_{kj} - X_{nj} \geq M(\hat{u}_{nkj} - 1) + \epsilon, \quad (4)$$
$$X_{nj} - l_{kj} \leq M\hat{l}_{nkj}, \quad (5)$$
$$X_{nj} - l_{kj} \geq M(\hat{l}_{nkj} - 1) + \epsilon, \quad (6)$$
$$u_{kj} \leq U_j \quad (7)$$
$$l_{kj} \geq L_j \quad (8)$$

A small number ϵ is used to force $\hat{u}_{nkj} = 1$ when $X_{nj} = u_{kj}$, and force $\hat{l}_{nkj} = 1$ when $X_{nj} = l_{kj}$. U_j and L_j denote the maximum and minimum value of attribute j. Constraints (7) and (8) bound on u_{kj} and l_{kj} to ensure a bounded M for computation efficiency.

It is possible that the upper and lower bounds apply to all examples, when both constraints (7) and (8) are binding. In that case, the literal does not have any classification power. We need numerical literals that are meaningful, or what we call *substantive*, that their ranges only apply to a subset of training examples. This means at most one of constraints (7) and (8) can be binding. Let a binary variable δ_{kj} indicates if literal j in pattern k is substantive. We use a

big M formulation to construct the constraints. Therefore $\forall k, \forall j \in \mathcal{J}_{\text{n}}$,

$$M\delta_{kj} \geq U_j - u_{kj}, \tag{9}$$
$$M\delta_{kj} \geq l_{kj} - L_j, \tag{10}$$
$$M(1-\delta_{kj}) \geq (u_{kj} - l_{kj}) - (U_j - L_j) + \epsilon, \tag{11}$$

Constraint (9) means $\delta_{kj} = 1$ if $u_{kj} < U_j$. Constraint (10) means $\delta_{kj} = 1$ if $l_{kj} > L_j$. (11) forces $\delta_{kj} = 0$ when $u_{kj} = U_j$ and $l_{kj} = L_j$, i.e., literal j is non-substantive.

2.1.2 Literals for Categorical Attributes

For a categorical attribute, a literal j has the form "$X_{\cdot j}$ = the v-th category", where $v \in \{1, ...V_j\}$ is an index for categories of attribute j and V_j is the total number of categories. We use $o_{kjv} \in \{0, 1\}$ to indicate whether the v-th category of attribute j is present in literal j of pattern k. To determine if X_n satisfies the condition in literal j, let $\hat{o}_{nkj} \in \{0, 1\}$ indicate if X_{nj} equals the category contained in this literal. We binary code X_{nj} into X_{njv} such that $X_{njv} = 1$ if X_{nj} takes the v-th category of attribute j. Therefore $\hat{o}_{nkj} = 1$ if and only if there exists $v \in \{1, ...V_j\}$ such that $o_{kjv} = 1$ and $X_{njv} = 1$. We formulate it as the following. For $\forall n, k, \forall j \in \mathcal{J}_{\text{c}}$,

$$\hat{o}_{nkj} \leq \sum_v X_{njv} o_{kjv}, \tag{12}$$
$$V_j \hat{o}_{nkj} \geq \sum_v X_{njv} o_{kjv}, \tag{13}$$
$$\sum_{v \in \{1,...,V_j\}} o_{kjv} \leq 1. \tag{14}$$

(14) ensures each categorical literal contains at most one value. This constraint forces a pattern to have a fixed form. If we remove the constraint and allow a literal to take multiple values, a pattern could have the following form: "$5 \leq X_1 \leq 20$ AND X_2 = red or blue." It depends on the application and users' preference as to whether leave this constraint in the MIP. The model will work the same without changing the rest of the formulation.

We define that a categorical literal j is *substantive*, if there exists some $v \in \{1, ..., V_j\}$ such that $o_{kjv} = 1$, indicated by $\delta_{kj} \in \{0, 1\}$. For $\forall k, \forall j \in \mathcal{J}_{\text{c}}$,

$$\delta_{kj} \leq \sum_{v \in \mathcal{V}_j} o_{kjv}, \tag{15}$$
$$V_j \delta_{kj} \geq \sum_{v \in \mathcal{V}_j} o_{kjv}. \tag{16}$$

2.1.3 Counting Classification Errors

Given the literals, X_n satisfies pattern k if and only if it satisfies every literal in the pattern. For categorical attributes, we consider both cases where a literal is substantive, i.e., $\delta_{kj} = 1$, and we need $\hat{o}_{nkj} = 1$; or non-substantive, i.e. $\delta_{kj} = 0$, and $\hat{o}_{nkj}$ is always 0 for all v. For numerical attributes, when the literal is substantive, the MIP needs to check if a data point satisfies both upper and lower bounds of the range, indicated by $\hat{u}_{nkj}$ and $\hat{l}_{nkj}$; when the literal is non-substantive, i.e., $u_{kj} = U_j$ and $l_{kj} = L_j$, then $\hat{u}_{nkj} = 1$ and $\hat{l}_{nkj} = 1$ for all X_n. Using $\omega_{nk} \in \{0, 1\}$ to indicate if X_n satisfies pattern k, the above conditions can be formulated below. $\forall n, k$,

$$\zeta_k + \sum_{j \in \mathcal{J}_{\text{n}}} \left(\hat{u}_{nkj} + \hat{l}_{nkj}\right) + \sum_{j \in \mathcal{J}_{\text{c}}} (\hat{o}_{nkj} + 1 - \delta_{kj}) - (2|\mathcal{J}_{\text{n}}| + |\mathcal{J}_{\text{c}}|) \leq \omega_{nk}, \tag{17}$$

$$(2|\mathcal{J}_{\text{n}}| + |\mathcal{J}_{\text{c}}| + 1)\omega_{nk} \leq \zeta_k + \sum_{j \in \mathcal{J}_{\text{n}}} \left(\hat{u}_{nkj} + \hat{l}_{nkj}\right) + \sum_{j \in \mathcal{J}_{\text{c}}} (\hat{o}_{nkj} + 1 - \delta_{kj}). \tag{18}$$

Let $\xi_n \in \{0, 1\}$ indicate if a classification error is made, which means either a positive data point does not satisfy any pattern, or a negative data point satisfies at least one pattern. In both cases $\xi_n = 1$. These two situations are captured by constraints (19) and (20).

$$\xi_n + \sum_k^K \omega_{nk} \geq 1, \forall n \in \mathcal{I}^+, \tag{19}$$
$$K\xi_n \geq \sum_k^K \omega_{nk}, \forall n \in \mathcal{I}^-, \tag{20}$$

where $\mathcal{I}^+$ denotes the set of indices for positive examples and $\mathcal{I}^-$ denotes the set of indices for negative examples. K is the upper bound on the number of patterns that we allow the solution to have. MIP creates this K "boxes" that will be filled up as it searches in the solution space.

When the MIP is formulated, we do not know how many out of the K "boxes" the MIP will use. Therefore we introduce binary variables $\zeta_k \in \{0, 1\}$ to indicate if pattern k is non-empty in the final pattern set, which means it contains at least one substantive literal. For $\forall k$,

$$J\zeta_k \geq \sum_j \delta_{kj}. \tag{21}$$

Since we are minimizing the total number of patterns in the objective, the constraint will always be binding.

2.1.4 The Objective

Now we reprsent the objective using decision variables introduced before. The MIP minimizes

$$\frac{1}{N}\sum_{n=1}^N \xi_n + C_1 \sum_k^K \sum_j^J \delta_{kj} + C_2 \sum_{k=1}^K \zeta_k$$

over variables ω_{nk},u_{kj}, l_{kj}, $\hat{u}_{nkj}$, $\hat{l}_{nkj}$, o_{kjv}, $\hat{o}_{nkj}$, δ_{kj}, δ_j, ξ_n, and ζ_k, such that they satisfy constraints (3) to (21).

The complexity of the MIP comes from three aspects, 1) choosing the upper and lower boundaries for ranges in numerical literals, and picking categories for categorical literals, 2) deciding for each example if it satisfies every literal and every pattern, and 3) deciding how many patterns are constructed from the K "boxes." There are in total $\mathcal{O}(NKJ)$ constraints and $\mathcal{O}(NKJ)$ decision variables for this MIP, though the full matrix of variables corresponding to the mixed integer programming formulation is sparse since most literals operate only on a small subset of the data. This formulation can be solved efficiently for small to medium sized datasets. As the size of the dataset grows, the computation gets complicated. We might need a faster method that operates in an approximate way on a much larger scale, presented below.

3 Optimized Or's of And's with Approximations

To speed up the learning process, we propose **Optimized Or's of And's with Approximations** (OOAx), that separates from the optimization process, the first two previously mentioned aspects of complexity. OOAx uses a *pre-mining then selecting* approach. It takes advantage of mature pattern mining techniques to generate a set of patterns. Then a secondary criteria is applied to further screen the rules to form a candidate pattern set. Finally, an integer linear program (ILP) searches within these patterns set for an optimal set. This method consists of following three steps, *pattern mining*, *pattern screening* and *pattern selecting*.

3.1 Pattern Mining

There are many frequently used pattern mining methods such as FP-growth (Han, Pei, and Yin, 2000), Apriori (Agrawal, Srikant, and others, 1994), Eclat (Zaki et al., 1997), etc. In our implementation, we use FP-growth in python (Borgelt, 2005) that takes the binary coded data, and user specified minimum support and maximum length, to generate patterns that satisfy the two requirements. The algorithm runs sufficiently fast (usually less than a second for thousands of observations). Since the FP-growth algorithm handles binarized data, we discretize the numerical attributes by manually selecting thresholds for bins. For instance, $X = 3.5$ can be transformed into $2 \leq X \leq 4$, etc. Note that there are other pattern mining techniques that handle real-valued variables.

3.2 Pattern Screening

In the pattern mining step, the number of generated patterns is usually overwhelming for even a medium size data set. For instance, for the sleep apnea data set (which we will discuss in detail in the experiment sections) of size 1192 patients and 112 binary coded attributes, if the maximum length is 3 and the minimum support is 5%, millions of patterns are generated. Ideally, we would like the candidate pattern set to contain thousands of patterns for computational convenience. Therefore, we use a secondary criteria to further screen the patterns.

$$\text{Score}(z) = \text{InfoGain}(S|z) - \gamma l_z. \tag{22}$$

This criteria considers the classification power of a pattern, measured by information gain $\text{InfoGain}(S|z)$, and the sparsity, measured by the length of the pattern l_z. Information gain of pattern z on data S is $\text{InfoGain}(S|z) = H(S) - H(S|z)$, where $H(S)$ is the entropy of S, written as $H(S) = -\sum_i P_i \log P_i$. $H(S|z)$ is the conditional entropy of S. Using this criteria, we select a set of candidate patterns $\mathcal{P}$ of size $K_{\mathcal{P}}$.

To represent the sparseness of each pattern, we create a binary matrix $\mathbf{P}$ of size $K_{\mathcal{P}} \times J$, where each row represents which attribute is present in a pattern. For instance, $P_{kj} = 1$ indicates that literal j is substantive in pattern k, and $P_{kj} = 0$ otherwise. We also need to determine for each example, which of the $K_{\mathcal{P}}$ patterns it satisfies. For a data set with N examples, we create a matrix $\mathbf{W}$ of size $N \times K_{\mathcal{P}}$, where the k-th element in the n-th row, ω_{nk}, indicates if the n-th observation satisfies pattern k. Both matrices are pre-computed before the final step.

3.3 Pattern Selecting

The previous two steps greatly reduce the computational load by feeding the final step with a set of high quality candidate patterns. Now our goal is only to select an optimal set A from the candidate set $\mathcal{P}$. We formulate an ILP using the same objective (2), and present it below.

$$\min_{\xi_n,\zeta_k} \frac{1}{N}\sum_{n=1}^{N}\xi_n + C_1\sum_{k}^{K_{\mathcal{P}}}\zeta_k l_k + C_2\sum_{k}^{K_{\mathcal{P}}}\zeta_k$$

such that

$$\xi_n + \sum_{k=1}^{K_{\mathcal{P}}}\omega_{nk}\zeta_k \geq 1, \forall n \in \mathcal{I}^+ \tag{23}$$

$$K\xi_n \geq \sum_{k=1}^{K_{\mathcal{P}}}\omega_{nk}\zeta_k, \forall n \in \mathcal{I}^- \tag{24}$$

$$\xi_n, \zeta_k \in \{0,1\}. \tag{25}$$

The length of pattern k, l_k, can be pre-computed by $l_k = \sum_{j=1}^{J} P_{kj}$. Constraint (23) means that an error occurs for a positive example if it does not satisfy any patterns that are selected. Constraint (24) means that an error occurs for a negative example if it satisfies at least one pattern that is selected. This ILP only involves $\mathcal{O}(N)$ constraints and

$\mathcal{O}(N) + \mathcal{O}(K_{\mathcal{P}})$ variables, which is much simpler than the MIP in an OOA framework.

The difference between OOA and OOAx method is that the latter avoids forming patterns in the optimization process, by handling it to other efficient off-the-shelf algorithms. Separating the mining step from the optimization problem renders more control to users over the quality and size of desired patterns. Users can manually modify the pre-mining and screening process by applying domain-specific minimum support, maximum length and secondary selection criteria.

4 Analysis on Patterns and OA Models

In this section, we discuss the quality of patterns in an OA classifier. Certain properties of the patterns improve computation complexity. We also show the VC dimension of OA models and compare OA classifiers with other discrete classifiers (decision trees and random forests). Due to the page limit, some proofs are provided in the supplementary material.

4.1 Bounds on Patterns

Define the *support set* of pattern z over data set S as

$$\mathcal{I}^S(z) = \{X | \mathbb{1}_z(X) = 1, X \in S\}, \tag{26}$$

and the *support* of pattern z over S as

$$\text{supp}^S(z) = |\mathcal{I}^S(z)|. \tag{27}$$

$\text{supp}^{S^+}(z)$ is called the *positive support* of z, which is the number of positive examples in $\mathcal{I}^S(z)$, and $\text{supp}^{S^-}(z)$ is called the *negative support* of z, which is the number of negative examples in $\mathcal{I}^S(z)$.

An OA classifier is essentially an ensemble of weaker classifiers, patterns. Including patterns with a low quality is expensive, and as we will prove, unnecessary. First we show in Theorem 1 that the optimal solution never includes a pattern with a high negative support.

Theorem 1 *Take an OA model with regularization parameters C_1 and C_2. The OA model is trained on a data set S, consisting of N examples, N^+ of which are positive examples. If $A^* \in \arg\min_A L(A)$, then for any $z \in A^*$, $supp^{S^-}(z) \leq N^+ - N(C_1 + C_2)$.*

This means after pattern mining, we can safely reduce the pattern space by disregarding patterns with a negative support above $N^+ - N(C_1 + C_2)$. Similarly, we can also prove that if a pattern has a low positive support, removing it always achieves a better objective. Let $A_{\setminus z}$ denote the pattern set with pattern z removed from A.

Theorem 2 *Take an OA model with regularization parameters C_1 and C_2. The OA model is trained on a data set S, consisting N examples. If $supp^{S^+}(z) \leq (C_1 + C_2)N$, then $L(A_{\setminus z}) \leq L(A)$.*

It means we need not bother mining rules of low positive support. Theorem 2 is a stronger statement than Theorem 1 since it provides a lower bound on positive support for patterns in all pattern sets and saying that removing a low supported pattern always improves the performance; while Theorem 2 only applies to optimal solutions.

With the above theoretical guarantees, we know it is safe to reduce the pattern space, by setting the minimum positive support to be $(C_1 + C_2)N$ when we pre-mine the patterns and throwing away patterns with negative support higher than $N^+ - N(C_1 + C_2)$ in the screening stage. This does not benefit an OOA framework as it directly forms patterns. But it provides strong computational motivation for pre-mining patterns in an OOAx framework.

The sparseness of a model is also associated with the number of patterns in an OA model. We prove in Theorem 3 that the number of patterns in an optimal pattern set is upper bounded.

Theorem 3 *Take an OA model with regularization parameters C_1 and C_2. The OA model is trained on a set of N examples, N^+ of which are positive examples. If $A^* \in \arg\min_A L(A)$, then $|A^*| \leq \frac{N^+/N}{C_1+C_2}$.*

This theorem is meaningful not only for showing the simplicity of the output, but also gives us a suggestion for K when we use the MIP in an OOA framework. Knowing that the optimal set can never be larger than $\frac{N^+/N}{C_1+C_2}$, we can safely set K to be $\frac{N^+/N}{C_1+C_2}$. The smaller K can be set, the better it is computationally for the MIP.

4.2 VC Dimension of an OA classifier

Let us consider the VC dimension of hypothesis classes representing pattern sets selected from a pre-mined set $\mathcal{P}$. There are some results for k-DNF (Ehrenfeucht et al., 1989) and monotone functions (that is, Boolean functions that can be represented without negated literals) (Procaccia and Rosenschein, 2006). Littlestone (1988) has shown that the class of k-term monotone l-DNF formulas (i.e., with monomials containing at most l variables) has VC dimension at least $lk\lfloor \log(\frac{n}{m}) \rfloor$, where $l \leq m \leq n$, and $k \leq \binom{m}{l}$. However, his theorem does not have the constraint that the patterns come from a fixed pattern set.

Let $\mathcal{S} = \mathbb{R}^J$ represent the complete set of all possible data that could be constructed from J attributes. To compute the VC dimension, we introduce the following definition.

Definition 1 *An **Efficient Set** of $\mathcal{P}$ is a set of patterns where the support set of each pattern is not a subset of the rest of the efficient set, i.e.,*

$$\mathcal{P}^E = \{z | z \in \mathcal{P}, \mathcal{I}^{\mathcal{S}}(z) \not\subset \mathcal{I}^{\mathcal{S}}(\mathcal{P}^E_{\setminus z})\}.$$

This means for any pattern z in $\mathcal{P}^E$, there exists data points that satisfy only z and none of the rest of the patterns in $\mathcal{P}^E$. We call the efficient set with the maximum number of patterns the **Maximum Efficient Set** of $\mathcal{P}$, denoted as $\mathcal{P}^E_{\text{max}}$. We claim that the VC dimension of OOAx learned from $\mathcal{P}$ depends on the size of $\mathcal{P}^E_{\text{max}}$, stated as the following.

Theorem 4 *The VC dimension of an OA classifier f built from $\mathcal{P}$ equals the size of the maximum efficient set of $\mathcal{P}$:*

$$VCdim(f) := |\mathcal{P}^E_{max}|.$$

Proof 1 First we prove that $\text{VCdim}(f) \geq |\mathcal{P}^E_{\text{max}}|$, which means there exists a set of $|\mathcal{P}^E_{\text{max}}|$ examples $X_1, ...X_{|\mathcal{P}^E_{\text{max}}|}$ that any labels $Y_1, ..., Y_{|\mathcal{P}^E_{\text{max}}|}$ can be realized by a classifier f built from $\mathcal{P}$. To construct this example set, we use the maximum efficient set $\mathcal{P}^E_{\text{max}}$. For any pattern z_i in $\mathcal{P}^E_{\text{max}}$, since $\mathcal{I}^{\mathcal{S}}(z_i) \not\subset \mathcal{I}^{\mathcal{S}}(\mathcal{P}^E_{\text{max}} \backslash z_i)$, there always exists a data point $X_i \in \mathcal{S}$ that satisfies only z_i, i.e.,

$$X_i \in \mathcal{I}^{\mathcal{S}}(\mathcal{P}^E_{\text{max}} \backslash z) - \mathcal{I}^{\mathcal{S}}(z),$$

for $i \in \{1, ..., |\mathcal{P}^E_{\text{max}}|\}$. Each X_i is covered by exactly one pattern in $\mathcal{P}^E_{\text{max}}$. These points can always be shattered since for any labels, we can from a pattern set $A = \{z_i | z_i \in \mathcal{P}^E_{\text{max}}, \text{s.t. } \mathbb{1}_{z_i}(X_i) = 1, Y_i = 1\}$. Therefore, all possible labels of $Y_1, ..., Y_N$ can be realized, which means that $\text{VCdim}(f) \geq |\mathcal{P}^E_{\text{max}}|$.

Then we show $\text{VCdim}(f) \leq |\mathcal{P}^E_{\text{max}}|$. We prove this by contradiction. Let $\mathcal{P}^E_{\text{max}}$ be the maximum efficient set of $\mathcal{P}$. Assume there exists a set of h examples $X_1, ...X_h$ where $h > |\mathcal{P}^E_{\text{max}}|$, and their labels $Y_1, ...Y_h$ can always be realized. Let $\mathbf{0}_{\backslash i}$ denote an all-zero vector of size h except a one at the i-th position. For $\mathbf{0}_{\backslash i}$ to be a realizable set of labels, there must exist a pattern z_i that satisfies $\mathbb{1}_{z_i}(X_i) = 1$ and $\mathbb{1}_{z_i}(X_j) = 0$ for $j \neq i$. This should be true for all $i \in \{1, ..., h\}$. Therefore, there must exist h such patterns that each of them covers a data point that only satisfies this pattern. According to definition 1, this is equivalent to declaring that these h patterns is an efficient set, and the size of the set is h, which is greater than $|\mathcal{P}^E_{\text{max}}|$. This contradicts the assumption that $\mathcal{P}^E_{\text{max}}$ is the maximum efficient set and should contain the largest number of patterns. Therefore, $\text{VCdim}(f) \leq |\mathcal{P}^E_{\text{max}}|$.

Thus, we conclude that the VC-dimension of a classifier f built from $\mathcal{P}$ is $|\mathcal{P}^E_{\text{max}}|$. (Learning an efficient set will be another topic that we do not discuss in this paper.)

4.3 Comparing with Other Discrete Classifiers

Like OA classifiers, decision trees and random forests also discretize the input space and assign each subspace with a label. We prove that for these models, there always exist equivalent OA classifiers. These theorems are simple, but may not be obvious to those who have not thought about it. We present the definition of two classifiers being *equivalent* below.

Definition 2 *Two classifiers f_1, f_2 are equivalent if for any input X, $f_1(X) = f_2(X)$.*

In a decision tree, the leaves divide up the input space into areas with different labels, which will be the predicted outcome for any data that ends up in that area. A path from the root to a leaf is a conjunction of literals, i.e., a pattern. See Figure 1 as an example. The decision tree ends up with

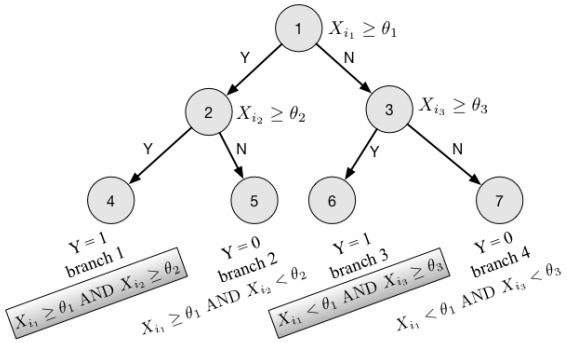

Figure 1: A decision tree and the corresponding patterns.

4 leaves, and therefore 4 patterns. To convert the tree into an equivalent OA classifier, we simply collect the patterns that are associated with positive leaves, in this case, leaf 4 and 6, shown in grey boxes.

Theorem 5 *For any decision tree f_1, there exists an equivalent OA classifier f_2, where the number of patterns in f_2 equals to the number of positive Y labels in f_1.*

Similar inductions hold for random forests. Random forest is an ensemble method based on decision trees. If an input data point falls into a positive leaf in at least half of all the trees, then it is labeled as positive. Therefore, the equivalent OA classifier consists of patterns that are conjunctions of positive patterns from at least half of all the trees. We summarize the above statements into Theorem 6.

Theorem 6 *For any random forest f_1, there exists an equivalent OA classifier f_2. If f_1 consists of K_{rf} decision trees, and the k-th tree has n_k positive leaves for $k \in \{1, ..., K_{rf}\}$, then the size of the pattern set in f_2 is upper bounded by $\sum_{\pi \in \Pi} \prod_{k \in \pi} n_k$, where Π is a collection of all possible combinations of $\lfloor \frac{K_{rf}}{2} \rfloor + 1$ elements selected from $\{1, ..., K\}$.*

The size is upper bounded by instead of exactly equal to $\sum_{\pi \in \Pi} \prod_{k \in \pi} n_k$ because some patterns could be equivalent, or contained in others. Note that we only need conjunction of $\lfloor \frac{K_{\text{rf}}}{2} \rfloor + 1$ patterns because conjunctions of more than that are contained in conjunctions of exactly $\lfloor \frac{K_{\text{rf}}}{2} \rfloor + 1$ positive patterns.

The above theorems provide theoretical guarantees that OA classifiers can be as good as decision trees and random forests, in terms of predictive performance, although it might not be desired to create complex OA models, since the whole purpose of designing an OA model is to favor its interpretability over other models.

5 Experiments

Our experiments include applying OA models to diagnose obstructive sleep apnea (OSA) and experimenting on 9 public datasets from UCI Machine Learning repository (Lichman, 2013).

To construct simple OA models for interpretability purposes, we set the maximum number of patterns to be 5 in all experiments. (In an OOA framework, we set $K = 5$; in a OOAx framework, we add a constraint that the sum of ζ_k's is less than or equal to 5.) Since we placed strong restrictions on the characteristics of OA models, we expect to lose predictive accuracy over unrestricted baseline methods. In many of the experiments we did, we found that OA models do not lose in performance, and most of the time are the top performing models, while achieving a substantial gain in interpretability.

5.1 Diagnosing Obstructive Sleep Apnea

The main experimental result is an application of OA models to build a diagnostic screening tool based on routinely available medical information. We analyzed polysomnography and self-reported clinical information from 1922 patients tested in the clinical sleep laboratory of Massachusetts General Hospital, first analyzed in 2015 (Ustun and Rudin, 2015; Ustun et al., 2015). The goal is to classify which patients who enter into the MGH Sleep Lab have OSA, based on a survey filled out upon admission. We produce predictive models for OSA screening using attributes from self-reported symptoms and self-reported medical information. The attributes include detailed information such as age, sex, BMI, sleepiness, if the patient snores, if the patient wakes up during sleep, if the patient falls back to sleep easily, the level of tiredness, etc. The data set was binary coded into 112 attributes.

Due to the size of this dataset, we chose OOAx for faster computation. We mined patterns with minimum support of 5% and maximum length of 3. We tuned parameters C_1 and C_2 using nested cross-validation to obtain the best performance, under the constraint that the pattern size cannot exceed 5. We measured out-of-sample performance using accuracy from 5-fold cross validation for OA models and 5 other methods that adhere to a certain level of interpretability, BOA (Wang et al., 2015), Lasso, C4.5, CART and RIPPER. For all baseline methods, we tuned the hyperparameters with grid search in nested cross validation. The results are displayed in Table 1.

Table 1: Accuracy comparison for OA models and baselines on obstructive sleep apnea dataset.

	Accuracy	Complexity
OOAx	.80(.01)	total number of patterns = 2.8 average length of patterns = 1.67 total number of literals = **4.7**
BOA	.80(.01)	total number of patterns = 4 average length of patterns = 1.75 total number of literals =**7**
RIPPER	.79(.01)	total number of rules = 4.6 average length of rules = 4 total number of literals = **18.4**
C4.5	.79(.01)	depth = 5 total number of nodes = **14.4**
CART	.79(.01)	depth = 5 total number of nodes = **11**
Lasso	.80(.01)	non-negative coefficients = **5**

To compare interpretability, we reported the complexity of each model averaged across 5 folds. For OA and BOA models, we reported the total number of patterns, average length of patterns and the total number of literals. OA models achieve the same performance as BOA models but with higher interpretability. This is due to a more flexible control over the size and shape of the pattern set compared to BOA models. RIPPER models are decision lists, having a different form than Or's of And's. We reported the total number of rules, average length of rules and total number of literals. For decision trees C4.5 and CART, we reported the depth of a tree, and the total number of nodes in a tree. For lasso, we reported the number of non-negative coefficients. Since baseline models have different logical forms than OA models, we compare one universal metric, the number of literals/nodes used in each model, marked in bold in Table 1. We find that OA models used substantially fewer literals than all other models while achieving a competitive accuracy to all models.

An example of an OA model is shown below.

if a patient satisfies (age $\geq$ 30 *AND* patient checked snoring as a potential symptom in the questionnaire),
OR (age $\geq$ 30 *AND* patient checked snoring as a reason for "why are you here" in the questionnaire),
OR (age $\geq$ 30 *AND* has hypertension),
OR (BMI $\geq$ 25) **then**
predict the patient has sleep apnea,
else
predict the patient does not have sleep apnea.
end if

The model lists four patterns to characterize patients that has sleep apnea. It is a sparse model with only a few attributes and a simple structure, and can potentially be used

Table 2: Accuracy comparison for OA models and baselines on UCI datasets.

	Data Type	*Interpretable Models*							*Uninterpretable Models*	
		OOA	OOAx	BOA	Lasso	C4.5	CART	RIPPER	random forest	SVM
blogger	Categorical	.85(.11)	**.86(.10)**	.80(.04)	.81(.08)	.77(.05)	.78(.07)	.76(.06)	.82(.07)	.82(.10)
votes		**.98(.02)**	**.98(.02)**	.95(.02)	.96(.02)	.96(.02)	.96(.03)	.96(.01)	.95(.02)	.97(.02)
tic-tac-toe		**1(.00)**	**1(.00)**	**1(.00)**	.71(.02)	.92(.03)	.93(.02)	.98(.01)	.99(.00)	.99(.00)
monks1		**1(.00)**	**1(.00)**	**1(.00)**	.76(.02)	.90(.06)	.88(.07)	.94(.12)	**1(.00)**	**1(.00)**
bupa	Numerical	.65(.02)	.65(.03)	.66(.02)	.68(.04)	.63(.04)	.68(.03)	.65(.05)	.70(.03)	**.73(.04)**
transfusion		.78(.02)	**.80(.01)**	.77(.01)	.77(.02)	.76(.02)	.78(.02)	.78(.02)	.78(.02)	.80(.02)
banknote		.98(.01)	.97(.01)	.96(.01)	1(.00)	.90(.01)	.90(.02)	.91(.01)	.91(.01)	**1(.00)**
indian-diabetes		.73(.02)	**.77(.03)**	.74(.02)	.67(.01)	.66(.03)	.67(.01)	.67(.02)	.76(.02)	.69(.01)
heart	mixed	.80(.04)	.84(.05)	.83(.07)	.85(.04)	.76(.06)	.77(.06)	.78(.04)	.81(.06)	**.86(.06)**

by people without a machine learning background.

5.2 Performance on UCI Datasets

We applied OOA and OOAx to several UCI datasets and compared with 5 previously mentioned interpretable models and 2 black box models, random forest and SVM. In the experimental set up, we set a time limit for the MIP in the OOA framework to ensure that it returns a solution in a reasonable amount of time. Table 2 displays the mean and standard deviation of out-of-sample accuracy across 5 folds.

We observed that even with the severe restrictions, OA classifiers achieve very competitive performance. For the four categorical datasets in Table 2, OA classifiers always do better than other models. Especially for tic-tac-toe and monks, where there are correct models that correctly classify all examples, OA models are able to discover the correct patterns and achieve 100% accuracy. For numerical and mixed datasets, OA models' performance levels are on par with those of other methods, sometimes slightly dominated by uninterpretable machine learning models.

We show an example of an OA classifier learned from dataset "votes" using OOA framework. This data set includes votes for each of the U.S. House of Representatives Congressmen on 16 key votes on water project cost sharing, duty free exports, immigration, education spending, anti-satellite test ban and etc. The objective is to predict if the voter is democratic or republican.

if a voter (votes for eduction spending *AND* for physician fee freeze *AND* against water project cost sharing),
OR (votes for export administration act of South Africa *AND* for physician fee freeze *AND* agains synfuels corporation cutback),
OR (votes against aid to Nicaraguan Contrast *AND* against adoption of the budget resolution *AND* against handicapped infants and toddlers act *AND* against superfund right to sue),
OR (votes for adoption of the budget resolution *AND* for physician fee freeze *AND* agains synfuels corporation cutback),
OR (votes against adoption of the budget resolution *AND* for El Salvador aid *AND* for physician fee freeze),
OR (votes for aid to Nicaraguan Contras *AND* against adoption of the budget resolution *AND* against duty free exports *AND* against synfuels corporation cutback), **then**
predict the voter is republican,
else
predict the voter is democratic.
end if

6 Conclusion

OA models have a long history. They are particularly useful as either (i) interpretable screening mechanisms, where they reduce much of the data from consideration from a further round of modeling, and (ii) consideration sets from marketing, which are rules that humans create to reduce cognitive load in order to make a decision.

We presented two optimization-based frameworks for learning Or's of And's. The first framework, OOA, uses a MIP to directly form patterns from data. It can deal with both categorical and numerical data without preprocessing. The second framework OOAx reduces computation through pre-mining patterns. We provided bounds on the support of patterns that guarantee that the pattern space can be safely reduced. Both methods can produce high quality OA classifiers, as demonstrated through experiments. They achieve competitive performance compared to other classifiers, with a substantial gain in sparsity and interpretability.

One of the main benefits not discussed extensively earlier is the benefit of customizability. Because we use MIP/ILP technology, constraints of almost any kind are very easy to include, and we do not need to derive a new algorithm; this benefit does not come with any other technology that we know of. Customizability is an important component of interpretability.

References

Agrawal, R.; Srikant, R.; et al. 1994. Fast algorithms for mining association rules. In *Proc. 20th Int. Conf. Very Large Data Bases (VLDB)*, volume 1215, 487–499.

Borgelt, C. 2005. An implementation of the fp-growth algorithm. In *Proc of the 1st interl workshop on open source data mining: frequent pattern mining implementations*, 1–5. ACM.

Chen, G.; Liu, H.; Yu, L.; Wei, Q.; and Zhang, X. 2006. A new approach to classification based on association rule mining. *Decision Support Systems* 42(2):674–689.

Cheng, H.; Yan, X.; Han, J.; and Hsu, C.-W. 2007. Discriminative frequent pattern analysis for effective classification. In *Data Engineering, 2007. ICDE 2007. IEEE 23rd International Conference on*, 716–725. IEEE.

Cohen, W. 1995. Fast effective rule induction. In *Proceedings of the twelfth international conference on machine learning*, 115–123.

Dawes, R. M. 1979. The robust beauty of improper linear models in decision making. *American psychologist* 34(7):571–582.

Ehrenfeucht, A.; Haussler, D.; Kearns, M.; and Valiant, L. 1989. A general lower bound on the number of examples needed for learning. *Information and Computation* 82(3):247–261.

Gilbride, T. J., and Allenby, G. M. 2004. A choice model with conjunctive, disjunctive, and compensatory screening rules. *Marketing Science* 23(3):391–406.

Han, J.; Pei, J.; and Yin, Y. 2000. Mining frequent patterns without candidate generation. In *ACM SIGMOD Record*, volume 29, 1–12. ACM.

Hauser, J. R.; Toubia, O.; Evgeniou, T.; Befurt, R.; and Dzyabura, D. 2010. Disjunctions of conjunctions, cognitive simplicity, and consideration sets. *Journal of Marketing Research* 47(3):485–496.

Johnson-Laird, P.; Khemlani, S. S.; and Goodwin, G. P. 2015. Logic, probability, and human reasoning. *Trends in cognitive sciences* 19(4):201–214.

Klivans, A. R., and Servedio, R. 2001. Learning dnf in time $2^{O(n^{1/3})}$. In *Proc of the 33rd Annual ACM Symp on Theory of Computing*, STOC '01, 258–265. New York, NY, USA: ACM.

Lavrac, N., and Dzeroski, S. 1994. Inductive logic programming. In *WLP*, 146–160. Springer.

Li, W.; Han, J.; and Pei, J. 2001. Cmar: Accurate and efficient classification based on multiple class-association rules. In *ICDM 2001*, 369–376. IEEE.

Lichman, M. 2013. UCI machine learning repository.

Littlestone, N. 1988. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine learning* 2(4):285–318.

Ma, B.; Liu, W.; and Hsu, Y. 1998. Integrating classification and association rule mining. In *Proceedings of the 4th International Conf on Knowledge Discovery and Data Mining (KDD)*.

Miller, G. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *The Psychological Review* 63:81–97.

Muggleton, S., and De Raedt, L. 1994. Inductive logic programming: Theory and methods. *The Journal of Logic Programming* 19:629–679.

Procaccia, A. D., and Rosenschein, J. S. 2006. Exact vc-dimension of monotone formulas. *Neural Information ProcessingâĂŤLetters and Reviews* 10(7):165–168.

Ustun, B., and Rudin, C. 2015. Supersparse linear integer models for optimized medical scoring systems. *Machine Learning*. Accepted.

Ustun, B.; Westover, M. B.; Rudin, C.; and Bianchi, M. T. 2015. Clinical prediction models for sleep apnea: The importance of medical history over symptoms. *Journal of Clinical Sleep Medicine*. Accepted.

Wang, T.; Rudin, C.; Doshi-Velez, F.; Liu, Y.; Klampfl, E.; and MacNeille, P. 2015. Or's of and's for interpretable classification, with application to context-aware recommender systems. *arXiv preprint arXiv:1504.07614*.

Yin, X., and Han, J. 2003. Cpar: Classification based on predictive association rules. In *SDM*, volume 3, 369–376. SIAM.

Zaki, M. J.; Parthasarathy, S.; Ogihara, M.; Li, W.; et al. 1997. New algorithms for fast discovery of association rules. In *KDD*, volume 97, 283–286.

SUPPLEMENTARY MATERIAL

Proof 2 *(Of Theorem 1) The objective function of the optimal solution A^* is*

$$\begin{aligned} L(A^*) = & \frac{N^+ - supp^{S^+}(A^*) + supp^{S^-}(A^*)}{N} + \\ & C_1 \#literals(A^*) + C_2 \#patterns(A^*) \\ \geq & \frac{N^+ - supp^{S^+}(A^*) + supp^{S^-}(A^*)}{N} + C_1 + C_2 \\ \geq & \frac{supp^{S^-}(A^*)}{N} + C_1 + C_2. \end{aligned}$$

These inequalities become tight when A^* was one pattern with one literal covers the whole positive class and some of the negative class. Let $\emptyset$ denote an empty set where there are no patterns and all the data points are classified as negative, so the total number of errors are the number of positive data, denoted as N^+. The objective function given $\emptyset$ is

$$L(\emptyset) = \frac{N^+}{N}.$$

Since $A^* \in \arg\min_A L(A)$, $L(A^*) \leq L(\emptyset)$, then

$$supp^{S^-}(A) \leq N^+ - N(C_1 + C_2)$$

Since $I^{S^-}(A) = \cup_{a \in A} I_a^{S^-}$, then $supp^{S^-}(z) \leq supp^{S^-}(A^*)$, thus

$$supp^{S^-}(z) \leq N^+ - N(C_1 + C_2)$$

Proof 3 *(Of Theorem 2) The worst case when pattern z is removed is when z is an accurate rule with confidence equal to 100%, i.e., all data points that satisfy pattern z are positive; and the points covered by z are not covered by any other pattern. Therefore once removing it, the number of errors increased by the positive support of z. On the other hand, removing z benefits the regularization terms, by decreasing the sum of pattern lengths by at least 1, and the number of patterns by 1. Then the objective function given $A_{\setminus z}$ obeys*

$$\begin{aligned} L(A_{\setminus z}) = & L(A) + \frac{\#error(A_{\setminus z}) - \#error(A)}{N} + \\ & C_1 \left(\#literals(A_{\setminus z}) - \#literals(A)\right) + \\ & C_2 \left(\#patterns(A_{\setminus z}) - \#patterns(A)\right) \\ \leq & L(A) + \frac{supp^{S^+}(z)}{N} - C_1 - C_2. \end{aligned}$$

In order to prove $L(A_{\setminus z}) \leq L(A)$, we need

$$L(A) + \frac{supp^{S^+}(z)}{N} - C_1 - C_2 \leq L(A),$$

i.e.,

$$supp^{S^+}(z) \leq (C_1 + C_2) N.$$

Proof 4 *(Of Theorem 3) Let $M^* = |A^*|$. A^* contains M^* patterns where each pattern has at least one literal. Therefore, the objective function given A^* is lower bounded by*

$$\begin{aligned} L(A^*) & \geq \frac{\#error(A^*)}{N} + C_1 M^* + C_2 M^* \\ & \geq M^* (C_1 + C_2). \end{aligned}$$

Since $A^* \in \arg\min_A L(A)$, $L(A^*) \leq L(\emptyset)$. That is

$$M^* (C_1 + C_2) \leq \frac{N^+}{N}.$$

Therefore

$$M^* \leq \frac{N^+/N}{C_1 + C_2}.$$